(12)

Qualitative Reasoning about the Physical World

Final Report
March, 1985 to December, 1991

ONR Contract N00014-85K-0225
Computer Science Division
Office of Naval Research

Kenneth D. Forbus
Department of Computer Science
University of Illinois at Urbana-Champaign

(Current address)
The Institute for the Learning Sciences
Northwestern University
1890 Maple Avenue
Evanston, Illinois, 60201

email: forbus@ils.nwu.edu
voice: (708) 491-7699

1 Introduction

This is the final report for the project. The next section summarizes our research results, broken down by area. Finally, we describe our efforts towards technology transfer.

2 Research summary by area

Qualitative physics is the subfield of artificial intelligence which investigates the kinds of representations and reasoning techniques which allow engineers, scientists, and just plain folks to understand and interact with the physical world. For example, Navy trainees learn to operate, maintain, and troubleshoot complex physical systems without needing advanced degrees in physics. Their commonsense knowledge is augmented by their technical training and new skills, in order to carry out their tasks. To make computers that are capable of the same kind of flexible reasoning and learning regarding the physical world requires understanding commonsense physics well enough to program it.

Qualitative physics is important scientifically because investigating particular kinds of reasoning can shed light on human cognition. It is also important technologically, with substantial potential economic benefits. Today's CAD systems share little of an engineer's knowledge, and thus cannot function as capable assistants. Training needs skyrocket as the complexity of systems grows, and intelligent tutoring systems are one of the few solutions on the horizon. A central goal of qualitative physics is to develop the principles of representation and reasoning which can support such applications.

The intellectual foundation of the project is Qualitative Process (QP) theory, which provides a formal description of qualitative knowledge and reasoning techniques which allow it to be applied to tasks. Unlike the typical expert-systems knowledge base, which focuses on capturing what is known about a specific system that is relevant to a given task (e.g., diagnosis of aircraft engines), a QP domain theory encodes fundamental principles of an engineering domain. The domain theory can be used to describe a variety of specific systems, and to support a variety of reasoning tasks. For example, in this project we showed using simple domain models that QP theory can support prediction, planning, and measurement interpretation.

One measure of the impact of this project is that in 1991, the journal *Artificial Intelligence* published a special issue, "Qualitative Reasoning about physical systems II," which highlighted current work in the subfield of qualitative physics. Of the ten technical papers which appeared in that special issue, three of them were the results of this project.

Our research can be divided into three areas: Qualitative modeling, spatial reasoning, and developing new reasoning techniques. We describe each in turn, highlighting the key developments, providing a list of papers produced by the project, and what the students involved with each part are doing now.

2.1 Qualitative modeling

Our work on qualitative modeling can be divided into two aspects: Improvements in the methodology of qualitative modeling itself, and the construction of richer, more powerful qualitative models. We summarize each in turn.

2.1.1 Compositional modeling

One of the principle contributions of this project was the creation of the *compositional modeling* methodology for organizing engineering knowledge. The key feature of this approach is the use of explicit *modeling assumptions* to control the level of detail in particular analyses. When checking if a new design is feasible, for instance, an engineer might ignore faults at first to see if the proposed system can work at all. Once it is established that the design might work, potential failure modes are worth considering. This modulation of detail during the course of an analysis is common, and capturing this ability is a crucial aspect of qualitative physics.

Compositional modeling has been widely adopted by the qualitative physics community. For example, Stanford's Knowledge Systems Lab and Kuiper's group at University of Texas at Austin are now using and extending our compositional modeling framework.

This work was done in collaboration with Brian Falkenhainer. Falkenhainer received his Ph.D., and is now a researcher at Xerox PARC.

1. Falkenhainer, B., and Forbus, K. "Setting up large-scale qualitative models", Proceedings of AAAI-88, August, 1988.

2. Falkenhainer, B. and Forbus, K. "Compositional Modeling: Finding the Right Model for the Job", *Artificial Intelligence*, **51** (1-3), October, 1991.

3. Falkenhainer, B. and Forbus, K. "Compositional modeling of physical systems", in *Recent Progress in Qualitative Physics*, Faltings, B. and Struss, P. (Eds.), MIT Press, 1992.

2.1.2 Qualitative models of thermodynamics

To build systems that can perform as robustly and flexibly as engineers do requires capturing the full spectrum of knowledge that engineers use. We assume that qualitative knowledge is an essential aspect of what an engineer knows, serving to organize and guide the use of other kinds of knowledge. A principle goal of this project is to explore the nature and use of qualitative models by developing broad-coverage models of specific domains of Navy interest, such as thermodynamics.

In this project we made several important steps in this direction. First, we made a preliminary exploration the role of qualitative knowledge in solving quantitative problems. Second, we developed substantial domain theory for engineering thermodynamics which

attempts to provide a reasonable level of detail for fundamental phenomena, such as flows and phase changes, that might be used in teaching the basic physics underlying propulsion plants. The kinds of components modeled include containers, fluid paths, heat paths, and pumps. We include models of liquids and gases, albeit assuming only a single working substance. The processes modeled include fluid flow, heat flow, boiling, condensation, and pumping, at different levels of detail. Third, we investigated the constraints imposed by the causal intuitions of engineers and scientists on the representation of equations. Fourth, we developed a new ontology for fluids which can support the kind of reasoning engineers do when understanding how the properties of a fluid change as it goes through a nozzle or how a counter-current heat exchanger works.

This work was done in collaboration with John Collins and Gordon Skorstad. Both are still working on their Ph.D.s, with Collins expected to finish this May and Skorstad expected to finish this September. Collins is now an assistant professor at Miami University.

1. Skorstad, G. and Forbus, K. "Qualitative and quantitative reasoning about thermodynamics", Proceedings of the Cognitive Science Society, August, 1989.

2. Collins, J. and Forbus, K. "Building qualitative models of thermodynamic processes," manuscript available via anonymous ftp from multivac.ils.nwu.edu, in `pub/PAPERS/fsthermo.ps.Z`.

3. Skorstad, G. "Finding stable causal interpretations of equations," in Faltings, B. and Struss, P. (Eds.) *Recent Advances in Qualitative Physics*, MIT Press, 1992.

4. Skorstad, G. "Towards a qualitative lagrangian theory of fluid flow," *Proceedings of AAAI92.*

2.2 Spatial Reasoning

Understanding many physical systems, such as internal combustion engines, involve reasoning both about dynamics and about space. Qualitative Process theory provides a solid framework for qualitative dynamics. In this project we developed a complementary framework for qualitative spatial reasoning, the *Metric Diagram/Place Vocabulary* model of spatial reasoning. We demonstrated the utility of this framework by creating a program that could analyze fixed-axis mechanical systems, such as mechanical clocks. In fact, the problem of qualitatively understanding a mechanical clock well enough to simulate it was a long-standing challenge in the field; our group was the first to do it, in 1988. The overall framework and most of the technical innovations we made in this work (e.g., the use of configuration space to reason about kinematics) are now widely used by others in the field. Our next step was to integrate our progress in dynamical reasoning with the results from the CLOCK project, so that we could analyze complex systems such as pumps and internal combustion engines.

The students who developed the CLOCK system were Paul Nielsen and Boi Faltings. Nielsen received his Ph.D. and is now a researcher at General Electric's corporate labs. Boi Faltings received his Ph.D. and is now an assistant professor at the Swiss Federal Institute of Technology. Hyeunkyeun Kim is expected to receive her Ph.D. this spring, and is looking for academic or industrial jobs in the midwestern United States.

1. Faltings, B. "A theory of qualitative kinematics in mechanisms", University of Illinois at Urbana-Champaign, Department of Computer Science Technical Report No. UIUCDCS-R-86-1274, May, 1986

2. Forbus, K., Nielsen, P., and Faltings, B. "Qualitative kinematics: A framework" Proceedings of IJCAI-87, *Milan, Italy, August, 1987.*

3. Faltings, B. "Qualitative place vocabularies for mechanisms in configuration space", University of Illinois Technical Report No. UIUCDCS-R-87-1360, July, 1987

4. Nielsen, P. "A qualitative approach to mechanical constraint", Proceedings of the National Conference on Artificial Intelligence, August 1988; also in *Qualitative Reasoning about Physical Systems*, Daniel S. Weld and Johan deKleer Eds., Morgan Kaufman Publishers, Inc., 1990

5. Nielsen, P. "A qualitative approach to rigid body mechanics", Ph. D. Thesis, University of Illinois, 1988

6. Nielsen, P. "The role of abstraction in place vocabularies", Proceedings of the Conference of the Cognitive Science Society, 1989

7. Faltings, B. "Qualitative kinematics in mechanisms", *Artificial Intelligence* **41**(1), June, 1990

8. "Qualitative reasoning about the geometry of fluid flow" *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, Cambridge, MA, July, 1990.

9. Forbus, K., Nielsen, P. and Faltings, B. "Qualitative Spatial Reasoning: The CLOCK Project", *Artificial Intelligence*, **51** (1-3), October, 1991.

10. "Qualitative kinematics of linkages", Technical Report No. UIUCDCS-R-90-1603, Department of Computer Science, University of Illinois at Urbana-Champaign, May, 1990.

11. Kim, H. "Augmenting qualitative simulation with global filtering," *Proceedings of CogSci92.*

2.3 Reasoning Techniques

A crucial part of this project was pushing the state of the art in qualitative reasoning, and in AI reasoning techniques more generally. Many of the ideas and experience in assumption-based truth maintenance were worked out using examples from this project. In turn, we developed, in collaboration with Johan de Kleer at Xerox PARC, new techniques for building focused ATMS-based problem solvers. We developed new techniques envisioning and reasoning about inequalities. We developed novel techniques for measurement interpretation, to support the construction of programs which can summarize what is happening in complex systems in easily understood terms. We developed techniques for planning, as a step towards making qualitative physics useful for robotics and reasoning about procedures.

Several students were involved in this part of the work. John Hogge received his Master's degree and is programming in industry. Dennis DeCoste is expected to receive his Ph.D. this spring and will be working at NASA's Jet Propulsion Laboratory.

1. Forbus, K. "Interpreting measurements of physical systems" Proceedings of AAAI-86, Philadelphia, PA., August, 1986

2. Forbus, K. "The Qualitative Process Engine", Technical Report No. UIUCDCS-R-86-1288, December, 1986.

3. Forbus, K. "Interpreting observations of physical systems", *IEEE Systems, Man, and Cybernetics*, Special issue on Causal and diagnostic reasoning, Vol SMC-17, Number 3, pp 350-359, May/June, 1987.

4. Forbus, K. "The logic of occurrence", Proceedings of IJCAI-87, Milan, Italy, August, 1987.

5. Hogge, J. "Compiling plan operators from domains expressed in Qualitative Process theory", Proceedings of AAAI-87, July, 1987.

6. Hogge, J. "The compilation of planning operators from Qualitative Process theory models", Technical Report No. UIUCDCS-R-87-1368, September, 1987

7. Hogge, J. "TPLAN: A temporal interval-based planner with novel extensions", Technical Report No. UIUCDCS-R-87-1367, September, 1987

8. Forbus, K. and de Kleer, J. "Focusing the ATMS", Proceedings of AAAI-88, August, 1988.

9. Forbus, K. "Introducing actions into qualitative simulation", Proceedings of IJCAI-89, August, 1989.

10. Forbus, K. "QPE: A study in assumption-based truth maintenance" *International Journal of Artificial Intelligence in Engineering*, 1989.

11. DeCoste, D. "Dynamic across-time measurement interpretation," *Artificial Intelligence*, **51**,1-3, pp 273–434.

3 Technology Transfer

We have made several of the research systems developed as part of this project publically available to interested researchers. This includes ATMoSphere and QPE, described next.

ATMoSphere is a pattern-directed inference engine that is organized around an assumption-based truth maintence system (ATMS). It supports several styles of problem-solving using the ATMS, including the *implied-by* strategy we developed to support reasoning in unbounded domains. ATMoSphere is the substrate for most of our other systems.

QPE is our qualitative simulator. It takes as input a domain theory and a specific system to be modeled. It creates a model of the system using the vocabulary of the domain model, and produces an envisionment which represents all possible qualitatively distinct behaviors the system might undergo and the transitions between them. It also supports the use of simplifying assumptions to control reasoning with large-scale models. QPE has been used as a component in systems for measurement interpretation, planning, and learning physical theories.

Since we distribute these systems via anonymous ftp, we do not have exact records of who has received them or to what use they have been put. Often the first time we hear about someone using QPE is when they send email asking for help in modeling something. In 1991, requests included Siemens (Munich), the Israel Institute of Technology, USC's AI group, and NASA's Jet Propulsion Laboratory. In addition, two japanese researchers (Gitchang Okuno at NTT and Yoshihiko Ohta at ICOT) are using QPE as an example for research on parallel algorithms.

In addition to distributing research systems, Forbus has written a textbook which includes in its coverage many of the reasoning techniques developed as part of this project. The book is called *Building Problem Solvers*, and is being written in collaboration with Johan de Kleer (of Xerox PARC). Associated with the book is a substantial body of code (over 340 pages of listings) which will be made publically available via ftp upon publication. In addition to being useful for advanced students, we have designed it to be useful for programmers in industry and government who want to learn how to build sophisticated AI inference systems. We believe this combination of programs and textbook will prove to be a very effective means of technology transfer. The book should be available from MIT Press this summer.